



US009460475B2

(12) **United States Patent**
Chrapko et al.

(10) **Patent No.:** **US 9,460,475 B2**
(45) **Date of Patent:** **Oct. 4, 2016**

(54) **DETERMINING CONNECTIVITY WITHIN A COMMUNITY**

(56) **References Cited**

(71) Applicants: **Evan V Chrapko**, Edmonton (CA);
Leo M. Chan, Edmonton (CA)

6,356,902 B1 3/2002 Tan et al.
6,509,898 B2 1/2003 Chi et al.
6,633,886 B1 10/2003 Chong

(72) Inventors: **Evan V Chrapko**, Edmonton (CA);
Leo M. Chan, Edmonton (CA)

(Continued)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

FOREIGN PATENT DOCUMENTS

CA 2600344 8/2006
CA 2775899 4/2011

(Continued)

(21) Appl. No.: **14/282,935**

(22) Filed: **May 20, 2014**

(65) **Prior Publication Data**

US 2014/0258160 A1 Sep. 11, 2014

OTHER PUBLICATIONS

Din et al. "Transitive closure and metric inequality of weighted graphs: detecting protein interaction modules using cliques" Int. J. Data Mining and Bioinformatics. vol. I. No. 2, 2006, pp. 162-177, Sep. 2006.*

(Continued)

Related U.S. Application Data

(63) Continuation of application No. 13/498,429, filed as application No. PCT/CA2010/001531 on Sep. 30, 2010, now Pat. No. 9,171,338.

(60) Provisional application No. 61/247,343, filed on Sep. 30, 2009.

(51) **Int. Cl.**
G06Q 50/00 (2012.01)
G06Q 30/02 (2012.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06Q 50/01** (2013.01); **G06Q 30/02** (2013.01); **H04L 45/124** (2013.01); **H04L 67/22** (2013.01); **H04L 67/36** (2013.01)

(58) **Field of Classification Search**
CPC H04L 12/26
USPC 370/216, 238, 248, 252, 395.21
See application file for complete search history.

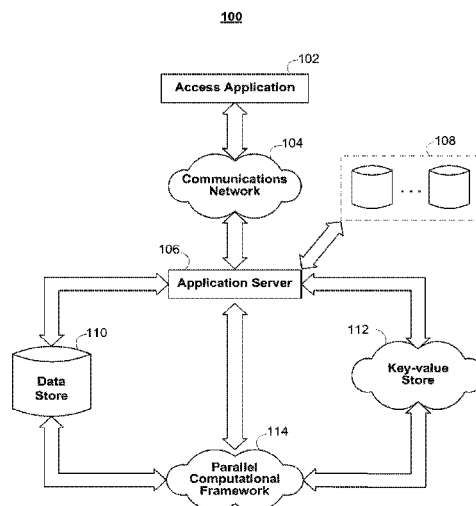
Primary Examiner — Dang Ton
Assistant Examiner — Nourali Mansoury

(74) *Attorney, Agent, or Firm* — Ropes & Gray LLP

(57) **ABSTRACT**

Systems and methods to determine trust scores and/or trustworthiness levels and/or connectivity are described herein. Trust and/or Trustworthiness and/or connectivity may be determined within, among or between entities and/or individuals. Social analytics and network calculations described herein may be based on user-assigned links or ratings and/or objective measures, such as data from third-party ratings agencies. The trust score may provide guidance about the trustworthiness, alignment, reputation, status, membership status and/or influence about an individual, entity, or group. The systems and methods described herein may be used to make prospective real-world decisions, such as whether or not to initiate a transaction or relationship with another person, or whether to grant a request for credit.

20 Claims, 10 Drawing Sheets



(51) **Int. Cl.**
H04L 12/721 (2013.01)
H04L 29/08 (2006.01)

2010/0076987 A1 3/2010 Schreiner
 2010/0309915 A1 12/2010 Pirbhai et al.
 2011/0246237 A1 10/2011 Vdovjak
 2012/0317149 A1 12/2012 Jagota et al.
 2013/0332740 A1 12/2013 Sauve et al.
 2014/0156274 A1 6/2014 You et al.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,738,777 B2 5/2004 Bliss et al.
 7,086,085 B1 8/2006 Brown et al.
 7,130,262 B1* 10/2006 Cortez et al. 370/216
 7,130,908 B1 10/2006 Pecus et al.
 7,266,649 B2 9/2007 Yoshida et al.
 7,451,365 B2 11/2008 Wang et al.
 7,512,612 B1 3/2009 Akella et al.
 7,664,802 B2* 2/2010 Aaltonen et al. 707/725
 7,743,208 B2 6/2010 Yoshida et al.
 8,108,536 B1 1/2012 Hernacki et al.
 8,156,558 B2 4/2012 Goldfeder et al.
 8,261,078 B2* 9/2012 Barriga et al. 713/171
 8,301,617 B2 10/2012 Muntz et al.
 8,392,590 B2 3/2013 Bouchard et al.
 8,572,129 B1 10/2013 Lee et al.
 8,601,025 B1 12/2013 Shajenko et al.
 8,683,423 B2 3/2014 Amaral et al.
 2003/0133411 A1 7/2003 Ise et al.
 2003/0227924 A1 12/2003 Kodialam et al.
 2004/0018518 A1 1/2004 Krieb et al.
 2004/0088147 A1 5/2004 Wang et al.
 2004/0122803 A1 6/2004 Dom et al.
 2005/0096987 A1 5/2005 Miyauchi
 2005/0243736 A1 11/2005 Faloutsos et al.
 2006/0271564 A1 11/2006 Meng Muntz et al.
 2007/0220146 A1 9/2007 Suzuki
 2008/0015916 A1 1/2008 Cossey et al.
 2008/0183378 A1 7/2008 Weidner
 2009/0024629 A1 1/2009 Miyauchi
 2009/0027392 A1 1/2009 Jadhav et al.
 2009/0043489 A1 2/2009 Weidner
 2009/0064293 A1 3/2009 Li et al.
 2009/0106822 A1 4/2009 Obasanjo et al.
 2009/0198562 A1 8/2009 Wiesinger et al.
 2009/0296568 A1 12/2009 Kitada
 2010/0004940 A1 1/2010 Choi et al.

FOREIGN PATENT DOCUMENTS

EP 1511232 A1 3/2005
 JP 2001-298453 A 10/2001
 JP 2003-259070 A 9/2003
 JP 2006-260099 A 9/2006
 WO WO-2009/109009 A1 9/2009

OTHER PUBLICATIONS

Chakraborty, S. et al., "TrustBAC-Integrating Trust Relationships into the RBAC Model for Access Control in Open Systems", Proceedings of the eleventh ACM symposium on Access Control Models and Technologies, SACMAT '06, pp. 49-58, Jun. 7-9, 2006 (Jul. 6, 2006).
 Gan, Z. et al., "A Novel Reputation Computing Model for Mobile Agent-Based E-Commerce Systems", Proceedings of the International Conference on Information Security and Assurance, ISA 2008, pp. 253-260, Apr. 24-26, 2008 (Apr. 24, 2008).
 Golbeck, J. et al., "Inferring Trust Relationships in Web-based Social Networks", Journal of ACM Transactions of Internet Technology (TOIT), vol. 6, issue 4, Nov. 2006.
 Huynh, T.D. et al., "An Integrated Trust and Reputation Model for Open Multi-Agent Systems", Journal of Autonomous Agents and Multi-Agent Systems, vol. 13, issue 2, Sep. 2006.
 Mui, L. et al., "A Computational Model of Trust and Reputation", Proceedings of the 35th Annual Hawaii International Conference on System Sciences, HICSS '02, vol. 7, pp. 2431-2439, Jan. 7-10, 2002 (Jul. 1, 2002).
 Safaei, M. et al., "Social Graph Generation & Forecasting Using Social network Mining", Proceedings of the 33rd Annual IEEE International Computer Software and Applications Conference, COMPSAC '09, pp. 31-35, Jul. 20-24, 2009 (Jul. 20, 2009).

* cited by examiner

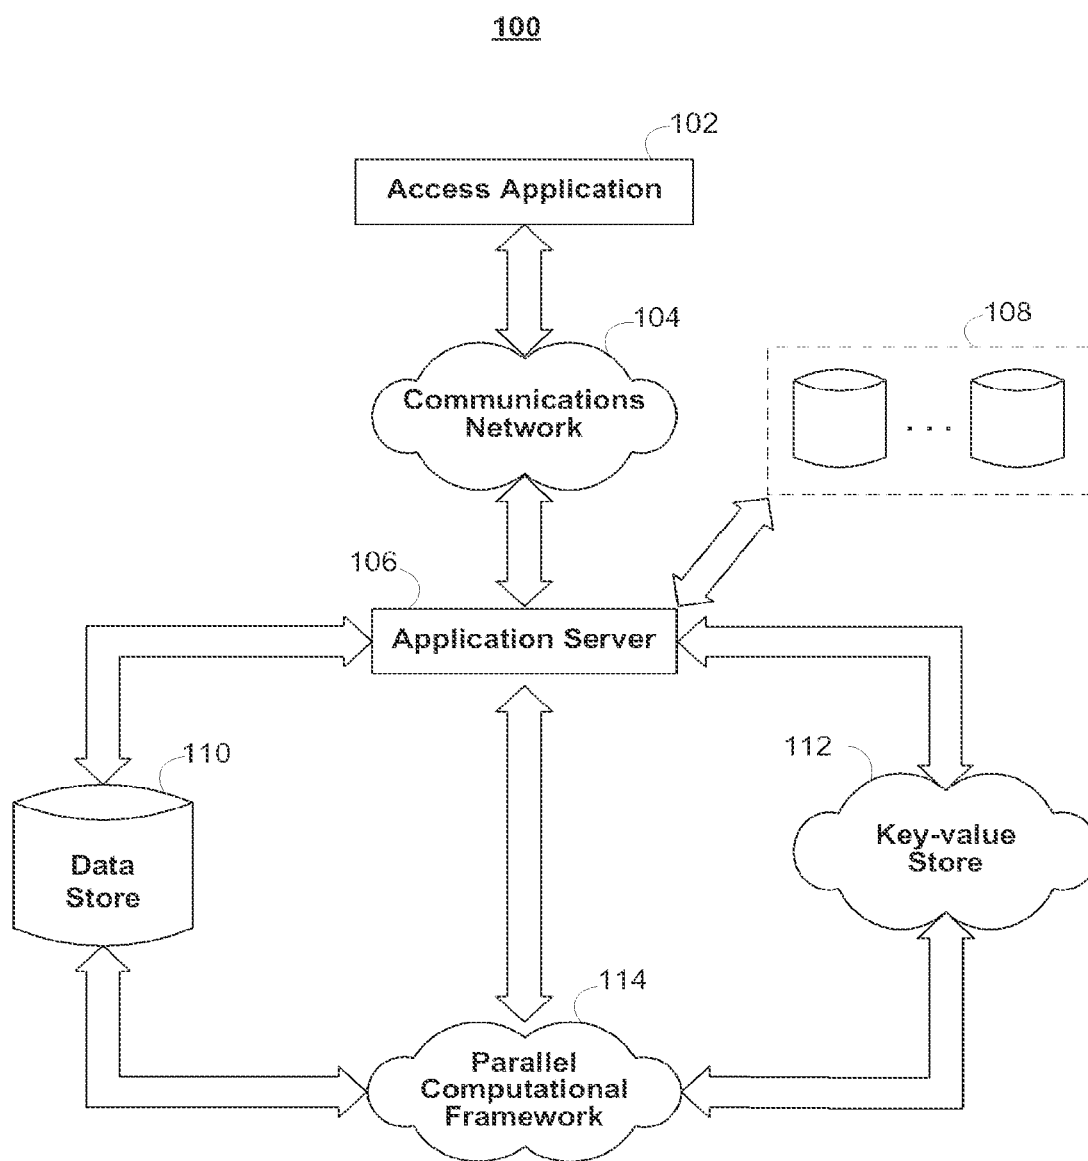


FIG. 1

200




FIG. 2

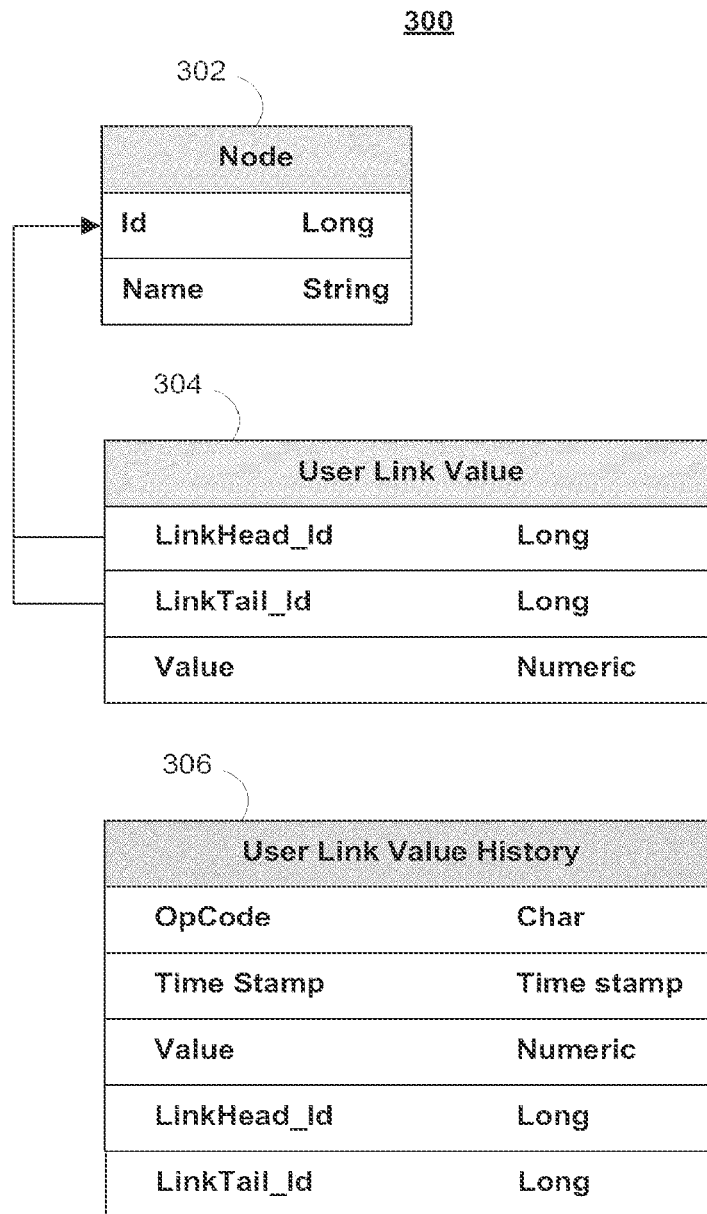


FIG. 3A

310

312

Node Table	
314	RowId 64-bit Integer
316	"info:inlinks" List
318	"info:outlinks" List
320	"info:oldHeadIdBuckets" List
322	"bucket:" + target node id List

FIG. 3B

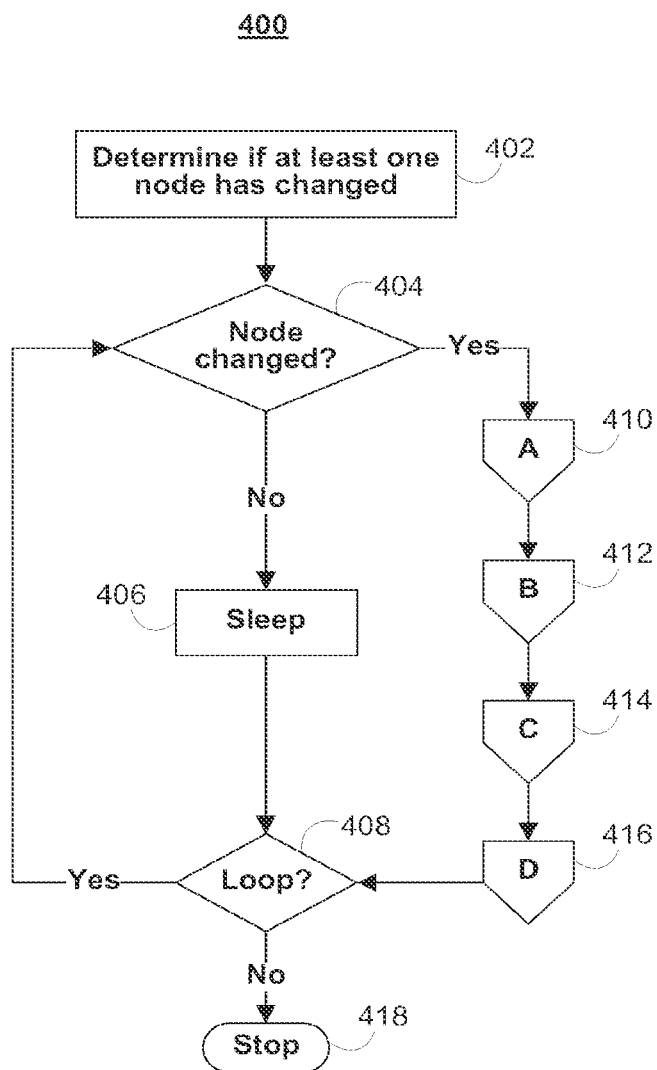


FIG. 4A

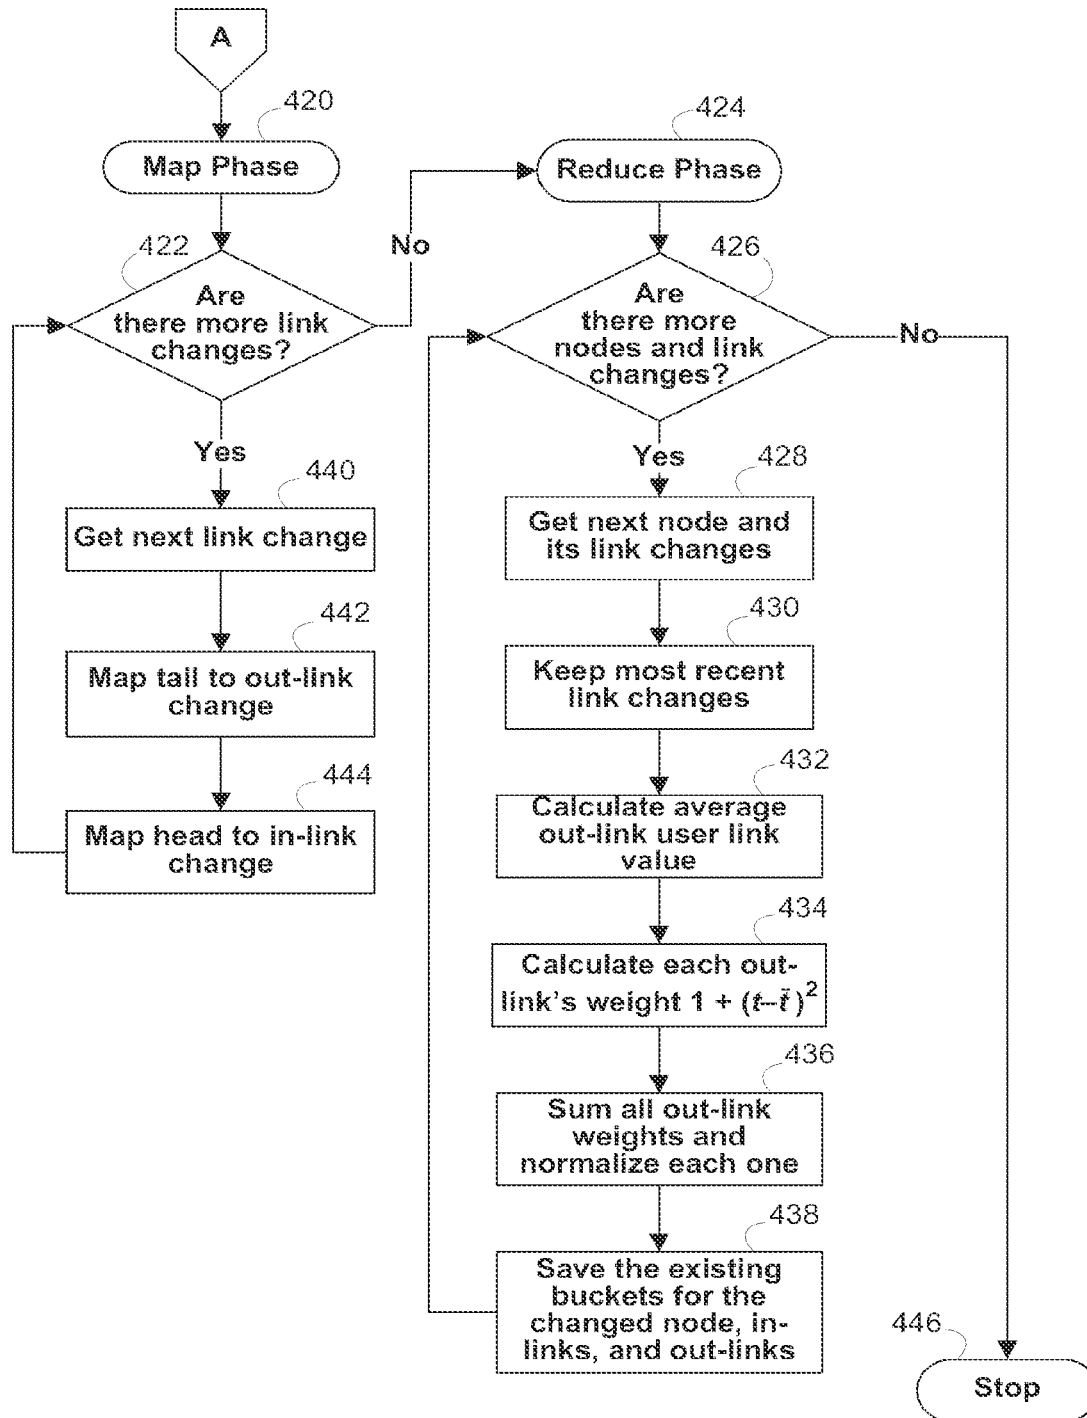


FIG. 4B

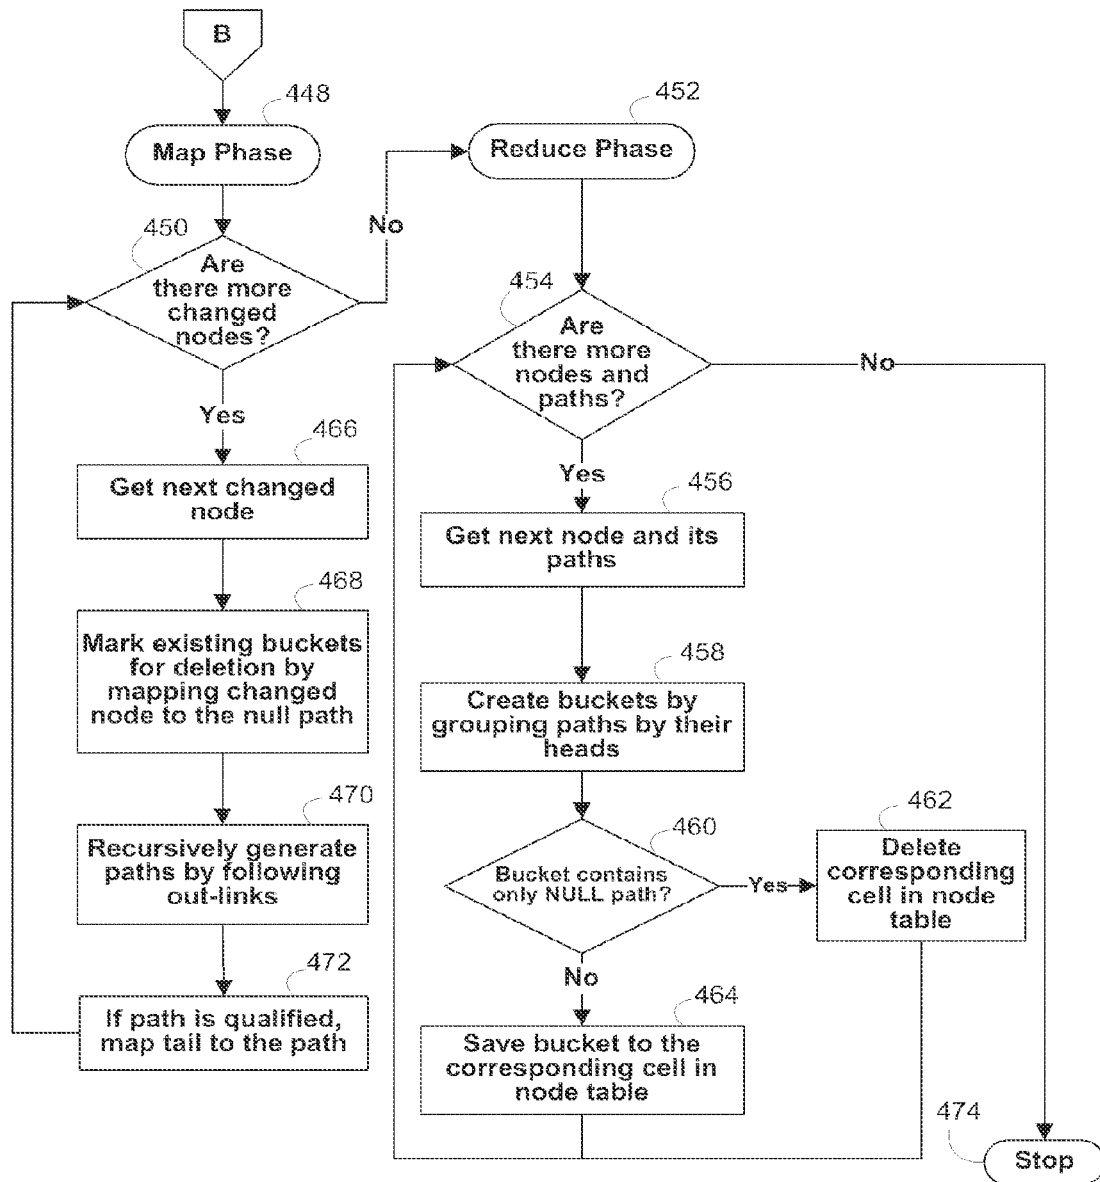


FIG. 4C

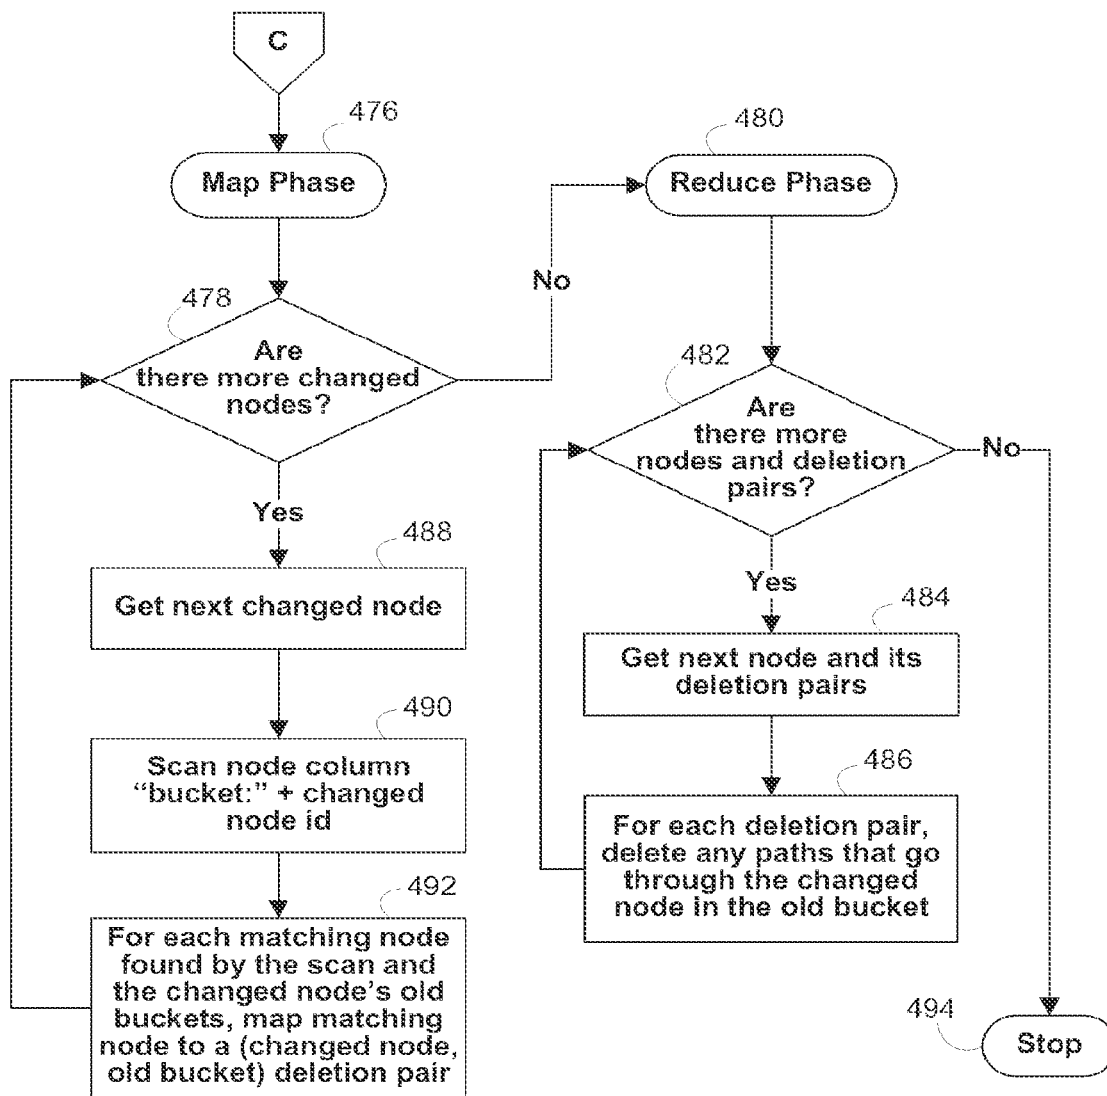


FIG. 4D

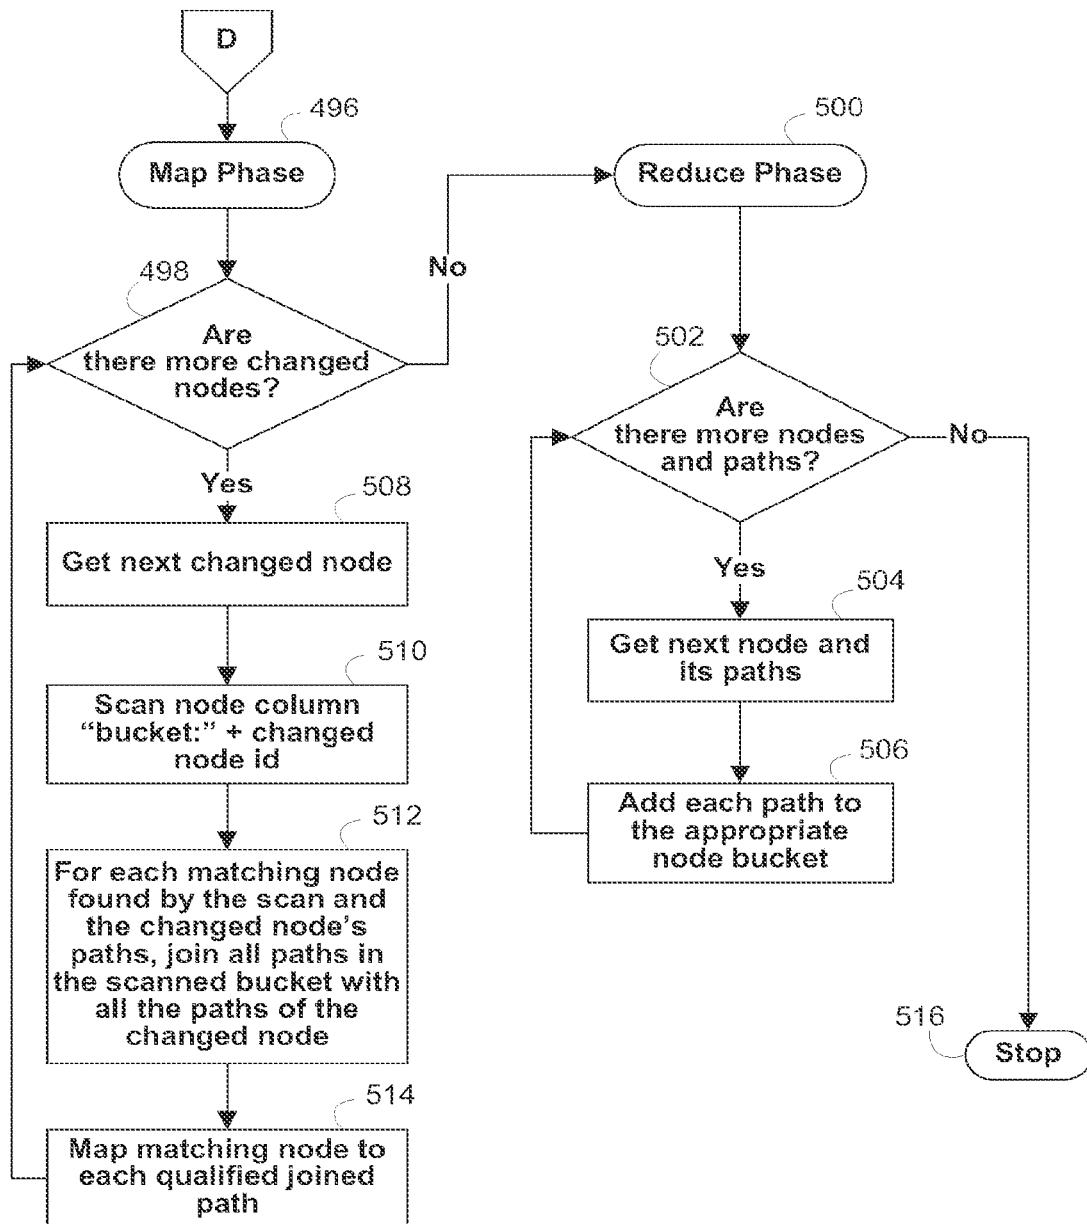


FIG. 4E

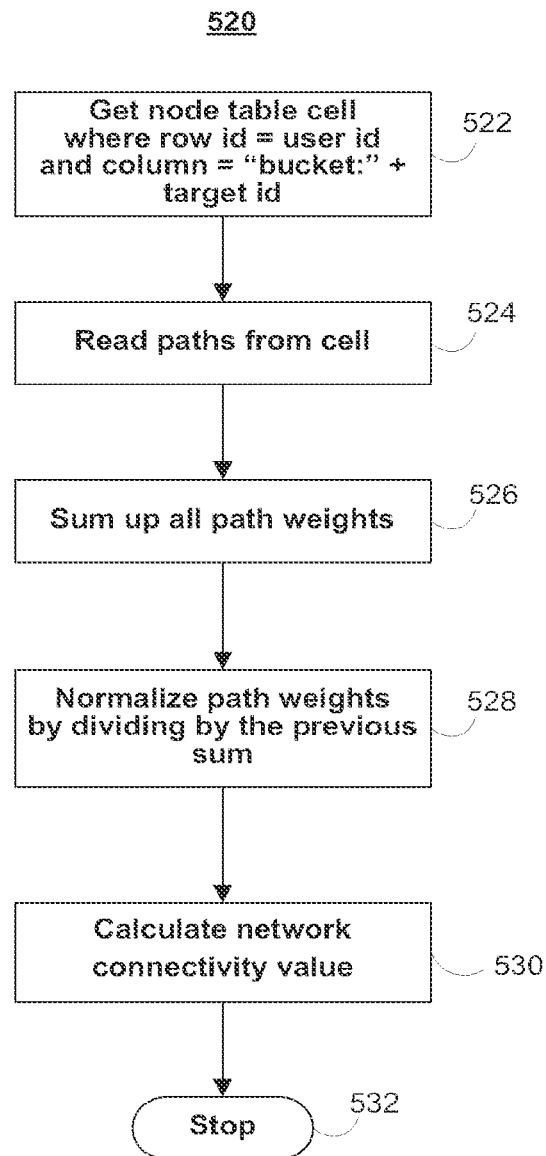


FIG. 5

1

DETERMINING CONNECTIVITY WITHIN A COMMUNITY

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation application of U.S. patent application Ser. No. 13/498,429, filed Mar. 27, 2012, which is a national stage filing under 35 U.S.C. § 371 of International Application No. PCT/CA2010/001531, filed Sep. 30, 2010, which claims the benefit of the filing date under 35 U.S.C. § 119(e) to U.S. provisional application No. 61/247,343, filed Sep. 30, 2009. The aforementioned, earlier-filed applications are hereby incorporated by reference herein in their entireties.

BACKGROUND OF THE INVENTION

This invention relates generally to networks of individuals and/or entities and network communities and, more particularly, to systems and methods for determining trust scores or connectivity within or between individuals and/or entities or networks of individuals and/or entities.

The connectivity, or relationships, of an individual or entity within a network community may be used to infer attributes of that individual or entity. For example, an individual or entity's connectivity within a network community may be used to determine the identity of the individual or entity (e.g., used to make decisions about identity claims and authentication), the trustworthiness or reputation of the individual or entity, or the membership, status, and/or influence of that individual or entity in a particular community or subset of a particular community.

An individual or entity's connectivity within a network community, however, is difficult to quantify. For example, network communities may include hundreds, thousands, millions, billions or more members. Each member may possess varying degrees of connectivity information about itself and possibly about other members of the community. Some of this information may be highly credible or objective, while other information may be less credible and subjective. In addition, connectivity information from community members may come in various forms and on various scales, making it difficult to meaningfully compare one member's "trustworthiness" or "competence" and connectivity information with another member's "trustworthiness" or "competence" and connectivity information. Also, many individuals may belong to multiple communities, further complicating the determination of a quantifiable representation of trust and connectivity within a network community. Even if a quantifiable representation of an individual's connectivity is determined, it is often difficult to use this representation in a meaningful way to make real-world decisions about the individual (e.g., whether or not to trust the individual).

Further, it may be useful for these real-world decisions to be made prospectively (i.e., in advance of an anticipated event). Such prospective analysis may be difficult as an individual or entity's connectivity within a network community may change rapidly as the connections between the individual or entity and others in the network community may change quantitatively or qualitatively. This analysis becomes increasingly complex as if applied across multiple communities.

SUMMARY OF THE INVENTION

In view of the foregoing, systems and methods are provided for determining the connectivity between nodes within

2

a network community and inferring attributes, such as trustworthiness or competence, from the connectivity. Connectivity may be determined, at least in part, using various graph traversal and normalization techniques described in more detail below.

In an embodiment, a path counting approach may be used where processing circuitry is configured to count the number of paths between a first node n_1 and a second node n_2 within a network community. A connectivity rating $R_{n_1 n_2}$ may then be assigned to the nodes. The assigned connectivity rating may be proportional to the number of subpaths, or relationships, connecting the two nodes, among other possible measures. Using the number of subpaths as a measure, a path with one or more intermediate nodes between the first node n_1 and the second node n_2 may be scaled by an appropriate number (e.g., the number of intermediate nodes) and this scaled number may be used to calculate the connectivity rating.

In some embodiments, weighted links are used in addition or as an alternative to the subpath counting approach. Processing circuitry may be configured to assign a relative user weight to each path connecting a first node n_1 and a second node n_2 within a network community. A user connectivity value may be assigned to each link. For example, a user or entity associated with node n_1 may assign user connectivity values for all outgoing paths from node n_1 . In some embodiments, the connectivity values assigned by the user or entity may be indicative of that user or entity's trust in the user or entity associated with node n_2 . The link values assigned by a particular user or entity may then be compared to each other to determine a relative user weight for each link.

The relative user weight for each link may be determined by first computing the average of all the user connectivity values assigned by that user (i.e., the out-link values). If t_i is the user connectivity value assigned to link i , then the relative user weight, w_i , assigned to that link may be given in accordance with:

$$w_i = 1 + (t_i - \bar{t})^2 \quad (1)$$

To determine the overall weight of a path, in some embodiments, the weights of all the links along the path may be multiplied together. The overall path weight may then be given in accordance with:

$$w_{path} = \prod(w_i) \quad (2)$$

The connectivity value for the path may then be defined as the minimum user connectivity value of all the links in the path multiplied by the overall path weight in accordance with:

$$t_{path} = w_{path} \times t_{min} \quad (3)$$

To determine path connectivity values, in some embodiments, a parallel computational framework or distributed computational framework (or both) may be used. For example, in one embodiment, a number of core processors implement an Apache Hadoop or Google MapReduce cluster. This cluster may perform some or all of the distributed computations in connection with determining new path link values and path weights.

The processing circuitry may identify a changed node within a network community. For example, a new outgoing link may be added, a link may be removed, or a user connectivity value may have been changed. In response to identifying a changed node, in some embodiments, the processing circuitry may recompute link, path, and weight

values associated with some or all nodes in the implicated network community or communities.

In some embodiments, only values associated with affected nodes in the network community are recomputed after a changed node is identified. If there exists at least one changed node in the network community, the changed node or nodes may first undergo a prepare process. The prepare process may include a “map” phase and “reduce” phase. In the map phase of the prepare process, the prepare process may be divided into smaller sub-processes which are then distributed to a core in the parallel computational framework cluster. For example, each node or link change (e.g., tail to out-link change and head to in-link change) may be mapped to a different core for parallel computation. In the reduce phase of the prepare process, each out-link’s weight may be determined in accordance with equation (1). Each of the out-link weights may then be normalized by the sum of the out-link weights (or any other suitable value). The node table may then be updated for each changed node, its in-links, and its out-links.

After the changed nodes have been prepared, the paths originating from each changed node may be calculated. Once again, a “map” and “reduce” phase of this process may be defined. During this process, in some embodiments, a depth-first search may be performed of the node digraph or node tree. All affected ancestor nodes may then be identified and their paths recalculated.

In some embodiments, to improve performance, paths may be grouped by the last node in the path. For example, all paths ending with node n_1 may be grouped together, all paths ending with node n_2 may be grouped together, and so on. These path groups may then be stored separately (e.g., in different columns of a single database table). In some embodiments, the path groups may be stored in columns of a key-value store implementing an HBase cluster (or any other compressed, high performance database system, such as BigTable).

In some embodiments, one or more threshold functions may be defined. The threshold function or functions may be used to determine the maximum number of links in a path that will be analyzed in a connectivity determination or connectivity computation. Threshold factors may also be defined for minimum link weights, path weights, or both. Weights falling below a user-defined or system-defined threshold may be ignored in a connectivity determination or connectivity computation, while only weights of sufficient magnitude may be considered.

In some embodiments, a user connectivity value may represent the degree of trust between a first node and a second node. In one embodiment, node n_1 may assign a user connectivity value of l_1 to a link between it and node n_2 . Node n_2 may also assign a user connectivity value of l_2 to a reverse link between it and node n_1 . The values of l_1 and l_2 may be at least partially subjective indications of the trustworthiness of the individual or entity associated with the node connected by the link. For example, one or more of the individual or entity’s reputation, status, and/or influence within the network community (or some other community), the individual or entity’s alignment with the trusting party (e.g., political, social, or religious alignment), past dealings with the individual or entity, and the individual or entity’s character and integrity (or any other relevant considerations) may be used to determine a partially subjective user connectivity value indicative of trust. A user (or other individual authorized by the node) may then assign this value to an outgoing link connecting the node to the individual or entity. Objective measures (e.g., data from third-party ratings agen-

cies or credit bureaus) may also be used, in some embodiments, to form composite user connectivity values indicative of trust. The subjective, objective, or both types of measures may be automatically harvested or manually inputted for analysis.

In some embodiments, a decision-making algorithm may access the connectivity values in order to make automatic decisions (e.g., automatic network-based decisions, such as authentication or identity requests) on behalf of a user. Connectivity values may additionally or alternatively be outputted to external systems and processes located at third-parties. The external systems and processes may be configured to automatically initiate a transaction (or take some particular course of action) based, at least in part, on received connectivity values. For example, electronic or online advertising may be targeted to subgroups of members of a network community based, at least in part, on network connectivity values.

In some embodiments, a decision-making algorithm may access the connectivity values to make decisions prospectively (e.g., before an anticipated event like a request for credit). Such decisions may be made at the request of a user, or as part of an automated process (e.g., a credit bureau’s periodic automated analysis of a database of customer information). This prospective analysis may allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other features of the present invention, its nature and various advantages will be more apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, and in which:

FIG. 1 is an illustrative block diagram of a network architecture used to support connectivity within a network community in accordance with one embodiment of the invention;

FIG. 2 is another illustrative block diagram of a network architecture used to support connectivity within a network community in accordance with one embodiment of the invention;

FIGS. 3A and 3B show illustrative data tables for supporting connectivity determinations within a network community in accordance with one embodiment of the invention;

FIGS. 4A-4E show illustrative processes for supporting connectivity determinations within a network community in accordance with one embodiment of the invention; and

FIG. 5 shows an illustrative process for querying all paths to a target node and computing a network connectivity value in accordance with one embodiment of the invention.

DETAILED DESCRIPTION

Systems and methods for determining the connectivity between nodes in a network community are provided. As defined herein, a “node” may include any user terminal, network device, computer, mobile device, access point, or any other electronic device. In some embodiments, a node may also represent an individual human being, entity (e.g., a legal entity, such as a public or private company, corporation, limited liability company (LLC), partnership, sole proprietorship, or charitable organization), concept (e.g., a social networking group), animal, or inanimate object (e.g., a car, aircraft, or tool). As also defined herein, a “network

5

community” may include a collection of nodes and may represent any group of devices, individuals, or entities.

For example, all or some subset of the users of a social networking website or social networking service (or any other type of website or service, such as an online gaming community) may make up a single network community. Each user may be represented by a node in the network community. As another example, all the subscribers to a particular newsgroup or distribution list may make up a single network community, where each individual subscriber may be represented by a node in the network community. Any particular node may belong in zero, one, or more than one network community, or a node may be banned from all, or a subset of, the community. To facilitate network community additions, deletions, and link changes, in some embodiments a network community may be represented by a directed graph, or digraph, weighted digraph, tree, or any other suitable data structure.

FIG. 1 shows illustrative network architecture 100 used to support the connectivity determinations within a network community. A user may utilize access application 102 to access application server 106 over communications network 104. For example, access application 102 may include a standard web browser, application server 106 may include a web server, and communication network 106 may include the Internet. Access application 102 may also include proprietary applications specifically developed for one or more platforms or devices. For example, access application 102 may include one or more instances of an Apple iOS, Android, or WebOS application or any suitable application for use in accessing application 106 over communications network 104. Multiple users may access application service 106 via one or more instances of access application 102. For example, a plurality of mobile devices may each have an instance of access application 102 running locally on the devices. One or more users may use an instance of access application 102 to interact with application server 106.

Communication network 104 may include any wired or wireless network, such as the Internet, WiMax, wide area cellular, or local area wireless network. Communication network 104 may also include personal area networks, such as Bluetooth and infrared networks. Communications on communications network 104 may be encrypted or otherwise secured using any suitable security or encryption protocol.

Application server 106, which may include any network server or virtual server, such as a file or web server, may access data sources 108 locally or over any suitable network connection. Application server 106 may also include processing circuitry (e.g., one or more microprocessors), memory (e.g., RAM, ROM, and hybrid types of memory), storage devices (e.g., hard drives, optical drives, and tape drives). The processing circuitry included in application server 106 may execute a server process for supporting the network connectivity determinations of the present invention, while access application 102 executes a corresponding client process. The processing circuitry included in application server 106 may also perform any of the calculations and computations described herein in connection with determining network connectivity. In some embodiments, a computer-readable medium with computer program logic recorded thereon is included within application server 106. The computer program logic may determine the connectivity between two or more nodes in a network community and it may or may not output such connectivity to a display screen or data

6

For example, application server 106 may access data sources 108 over the Internet, a secured private LAN, or any other communications network. Data sources 108 may include one or more third-party data sources, such as data from third-party social networking services and third-party ratings bureaus. For example, data sources 108 may include user and relationship data (e.g., “friend” or “follower” data) from one or more of Facebook, MySpace, openSocial, Friendster, Bebo, hi5, Orkut, PerfSpot, Yahoo! 360, LinkedIn, Twitter, Google Buzz, Really Simple Syndication readers or any other social networking website or information service. Data sources 108 may also include data stores and databases local to application server 106 containing relationship information about users accessing application server 106 via access application 102 (e.g., databases of addresses, legal records, transportation passenger lists, gambling patterns, political and/or charity donations, political affiliations, vehicle license plate or identification numbers, universal product codes, news articles, business listings, and hospital or university affiliations).

Application server 106 may be in communication with one or more of data store 110, key-value store 112, and parallel computational framework 114. Data store 110, which may include any relational database management system (RDBMS), file server, or storage system, may store information relating to one or more network communities. For example, one or more of data tables 300 (FIG. 3A) may be stored on data store 110. Data store 110 may store identity information about users and entities in the network community, an identification of the nodes in the network community, user link and path weights, user configuration settings, system configuration settings, and/or any other suitable information. There may be one instance of data store 110 per network community, or data store 110 may store information relating to a plural number of network communities. For example, data store 110 may include one database per network community, or one database may store information about all available network communities (e.g., information about one network community per database table).

Parallel computational framework 114, which may include any parallel or distributed computational framework or cluster, may be configured to divide computational jobs into smaller jobs to be performed simultaneously, in a distributed fashion, or both. For example, parallel computational framework 114 may support data-intensive distributed applications by implementing a map/reduce computational paradigm where the applications may be divided into a plurality of small fragments of work, each of which may be executed or re-executed on any core processor in a cluster of cores. A suitable example of parallel computational framework 114 includes an Apache Hadoop cluster.

Parallel computational framework 114 may interface with key-value store 112, which also may take the form of a cluster of cores. Key-value store 112 may hold sets of key-value pairs for use with the map/reduce computational paradigm implemented by parallel computational framework 114. For example, parallel computational framework 114 may express a large distributed computation as a sequence of distributed operations on data sets of key-value pairs. User-defined map/reduce jobs may be executed across a plurality of nodes in the cluster. The processing and computations described herein may be performed, at least in part, by any type of processor or combination of processors. For example, various types of quantum processors (e.g., solid-state quantum processors and light-based quantum

processors), artificial neural networks, and the like may be used to perform massively parallel computing and processing.

In some embodiments, parallel computational framework **114** may support two distinct phases, a “map” phase and a “reduce” phase. The input to the computation may include a data set of key-value pairs stored at key-value store **112**. In the map phase, parallel computational framework **114** may split, or divide, the input data set into a large number of fragments and assign each fragment to a map task. Parallel computational framework **114** may also distribute the map tasks across the cluster of nodes on which it operates. Each map task may consume key-value pairs from its assigned fragment and produce a set of intermediate key-value pairs. For each input key-value pair, the map task may invoke a user defined map function that transmutes the input into a different key-value pair. Following the map phase, parallel computational framework **114** may sort the intermediate data set by key and produce a collection of tuples so that all the values associated with a particular key appear together. Parallel computational framework **114** may also partition the collection of tuples into a number of fragments equal to the number of reduce tasks.

In the reduce phase, each reduce task may consume the fragment of tuples assigned to it. For each such tuple, the reduce task may invoke a user-defined reduce function that transmutes the tuple into an output key-value pair. Parallel computational framework **114** may then distribute the many reduce tasks across the cluster of nodes and provide the appropriate fragment of intermediate data to each reduce task.

Tasks in each phase may be executed in a fault-tolerant manner, so that if one or more nodes fail during a computation the tasks assigned to such failed nodes may be redistributed across the remaining nodes. This behavior may allow for load balancing and for failed tasks to be re-executed with low runtime overhead.

Key-value store **112** may implement any distributed file system capable of storing large files reliably. For example key-value store **112** may implement Hadoop’s own distributed file system (DFS) or a more scalable column-oriented distributed database, such as HBase. Such file systems or databases may include BigTable-like capabilities, such as support for an arbitrary number of table columns.

Although FIG. 1, in order to not over-complicate the drawing, only shows a single instance of access application **102**, communications network **104**, application server **106**, data source **108**, data store **110**, key-value store **112**, and parallel computational framework **114**, in practice network architecture **100** may include multiple instances of one or more of the foregoing components. In addition, key-value store **112** and parallel computational framework **114** may also be removed, in some embodiments. As shown in network architecture **200** of FIG. 2, the parallel or distributed computations carried out by key-value store **112** and/or parallel computational framework **114** may be additionally or alternatively performed by a cluster of mobile devices **202** instead of stationary cores. In some embodiments, cluster of mobile devices **202**, key-value store **112**, and parallel computational framework **114** are all present in the network architecture. Certain application processes and computations may be performed by cluster of mobile devices **202** and certain other application processes and computations may be performed by key-value store **112** and parallel computational framework **114**. In addition, in some

tations. For example, specially-configured routers or satellites may include processing circuitry adapted to carry out some or all of the application processes and computations described herein.

Cluster of mobile devices **202** may include one or more mobile devices, such as PDAs, cellular telephones, mobile computers, or any other mobile computing device. Cluster of mobile devices **202** may also include any appliance (e.g., audio/video systems, microwaves, refrigerators, food processors) containing a microprocessor (e.g., with spare processing time), storage, or both. Application server **106** may instruct devices within cluster of mobile devices **202** to perform computation, storage, or both in a similar fashion as would have been distributed to multiple fixed cores by parallel computational framework **114** and the map/reduce computational paradigm. Each device in cluster of mobile devices **202** may perform a discrete computational job, storage job, or both. Application server **106** may combine the results of each distributed job and return a final result of the computation.

FIG. 3A shows illustrative data tables **300** used to support the connectivity determinations of the present invention. One or more of tables **300** may be stored in, for example, a relational database in data store **110** (FIG. 1). Table **302** may store an identification of all the nodes registered in the network community. A unique identifier may be assigned to each node and stored in table **302**. In addition, a string name may be associated with each node and stored in table **302**. As described above, in some embodiments, nodes may represent individuals or entities, in which case the string name may include the individual or person’s first and/or last name, nickname, handle, or entity name.

Table **304** may store user connectivity values. In some embodiments, user connectivity values may be assigned automatically by the system (e.g., by application server **106** (FIG. 1)). For example, application server **106** (FIG. 1) may monitor all electronic interaction (e.g., electronic communication, electronic transactions, or both) between members of a network community. In some embodiments, a default user connectivity value (e.g., the link value 1) may be assigned initially to all links in the network community. After electronic interaction is identified between two or more nodes in the network community, user connectivity values may be adjusted upwards or downwards depending on the type of interaction between the nodes and the result of the interaction. For example, each simple email exchange between two nodes may automatically increase or decrease the user connectivity values connecting those two nodes by a fixed amount. More complicated interactions (e.g., product or service sales or inquires) between two nodes may increase or decrease the user connectivity values connecting those two nodes by some larger fixed amount. In some embodiments, user connectivity values between two nodes may always be increased unless a user or node indicates that the interaction was unfavorable, not successfully completed, or otherwise adverse. For example, a transaction may not have been timely executed or an email exchange may have been particularly displeasing. Adverse interactions may automatically decrease user connectivity values while all other interactions may increase user connectivity values (or have no effect). In addition, user connectivity values may be automatically harvested using outside sources. For example, third-party data sources (such as ratings agencies and credit bureaus) may be automatically queried for connectivity information. This connectivity information may include completely objective information, completely subjective information, composite information that is partially objec-

tive and partially subjective, any other suitable connectivity information, or any combination of the foregoing.

In some embodiments, user connectivity values may be manually assigned by members of the network community. These values may represent, for example, the degree or level of trust between two users or nodes or one node's assessment of another node's competence in some endeavor. As described above, user connectivity values may include a subjective component and an objective component in some embodiments. The subjective component may include a trustworthiness "score" indicative of how trustworthy a first user or node finds a second user, node, community, or subcommunity. This score or value may be entirely subjective and based on interactions between the two users, nodes, or communities. A composite user connectivity value including subjective and objective components may also be used. For example, third-party information may be consulted to form an objective component based on, for example, the number of consumer complaints, credit score, socio-economic factors (e.g., age, income, political or religious affiliations, and criminal history), or number of citations/hits in the media or in search engine searches. Third-party information may be accessed using communications network 104 (FIG. 1). For example, a third-party credit bureau's database may be polled or a personal biography and background information, including criminal history information, may be accessed from a third-party database or data source (e.g., as part of data sources 108 (FIG. 1) or a separate data source) or input directly by a node, user, or system administrator.

Table 304 may store an identification of a link head, link tail, and user connectivity value for the link. Links may or may not be bidirectional. For example, a user connectivity value from node n_1 to node n_2 may be different (and completely separate) than a link from node n_2 to node n_1 . Especially in the trust context described above, each user can assign his or her own user connectivity value to a link (i.e., two users need not trust each other an equal amount in some embodiments).

Table 306 may store an audit log of table 304. Table 306 may be analyzed to determine which nodes or links have changed in the network community. In some embodiments, a database trigger is used to automatically insert an audit record into table 306 whenever a change of the data in table 304 is detected. For example, a new link may be created, a link may be removed, or a user connectivity value may be changed. This audit log may allow for decisions related to connectivity values to be made prospectively (i.e., before an anticipated event). Such decisions may be made at the request of a user, or as part of an automated process, such as the processes described below with respect to FIG. 5. This prospective analysis may allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner. After such a change is detected, the trigger may automatically create a new row in table 306. Table 306 may store an identification of the changed node, and identification of the changed link head, changed link tail, and the user connectivity value to be assigned to the changed link. Table 306 may also store a timestamp indicative of the time of the change and an operation code. In some embodiments, operation codes may include "insert," "update," or "delete" operations, corresponding to whether a link was inserted, a user connectivity value was changed, or a link was deleted, respectively. Other operation codes may be used in other embodiments.

FIG. 3B shows illustrative data structure 310 used to support the connectivity determinations of the present invention. In some embodiments, data structure 310 may be stored

using key-value store 112 (FIG. 1), while tables 300 are stored in data store 110 (FIG. 1). As described above, key-value store 112 (FIG. 1) may implement an HBase storage system and include BigTable support. Like a traditional relational database management system, the data shown in FIG. 3B may be stored in tables. However, the BigTable support may allow for an arbitrary number of columns in each table, whereas traditional relational database management systems may require a fixed number of columns.

Data structure 310 may include node table 312. In the example shown in FIG. 3B, node table 312 includes several columns. Node table 312 may include row identifier column 314, which may store 64-bit, 128-bit, 256-bit, 512-bit, or 1024-bit integers and may be used to uniquely identify each row (e.g., each node) in node table 312. Column 316 may include a list of all the incoming links for the current node. Column 318 may include a list of all the outgoing links for the current node. Column 320 may include a list of node identifiers to which the current node is connected. A first node may be connected to a second node if outgoing links may be followed to reach the second node. For example, for $A \rightarrow B$, A is connected to B, but B may not be connected to A. As described in more detail below, column 320 may be used during the portion of process 400 (FIG. 4A) shown in FIG. 4B. Node table 312 may also include one or more "bucket" columns 322. These columns may store a list of paths that connect the current node to a target node. As described above, grouping paths by the last node in the path (e.g., the target node) may facilitate connectivity computations. As shown in FIG. 3B, in some embodiments, to facilitate scanning, bucket column names may include the target node identifier appended to the end of the "bucket:" column.

FIGS. 4A-4D show illustrative processes for determining the connectivity of nodes within a network community. FIG. 4A shows process 400 for updating a connectivity graph (or any other suitable data structure) associated with a network community. As described above, in some embodiments, each network community is associated with its own connectivity graph, digraph, tree, or other suitable data structure. In other embodiments, a plurality of network communities may share one or more connectivity graphs (or other data structure).

In some embodiments, the processes described with respect to FIG. 4A-4D may be executed to make decisions prospectively (i.e., before an anticipated event). Such decisions may be made at the request of a user, or as part of an automated process, such as the processes described below with respect to FIG. 5. This prospective analysis may allow for the initiation of a transaction (or taking of some particular action) in a fluid and/or dynamic manner.

At step 402, a determination is made whether at least one node has changed in the network community. As described above, an audit record may be inserted into table 306 (FIG. 3) after a node has changed. By analyzing table 306 (FIG. 3), a determination may be made (e.g., by application server 106 of FIG. 1) that a new link has been added, an existing link has been removed, or a user connectivity value has changed. If, at step 404, it is determined that a node has changed, then process 400 continues to step 410 (shown in FIG. 4B) to prepare the changed nodes, step 412 (shown in FIG. 4C) to calculate paths originating from the changed nodes, step 414 (shown in FIG. 4D) to remove paths that go through a changed node, and step 416 (shown in FIG. 4E) to calculate paths that go through a changed node. It should be noted that more than one step or task shown in FIGS. 4B,

11

4C, 4D, and 4E may be performed in parallel using, for example, a cluster of cores. For example, multiple steps or tasks shown in FIG. 4B may be executed in parallel or in a distributed fashion, then multiple steps or tasks shown in FIG. 4C may be executed in parallel or in a distributed fashion, then multiple steps or tasks shown in FIG. 4D may be executed in parallel or in a distributed fashion, and then multiple steps or tasks shown in FIG. 4E may be executed in parallel or in a distributed fashion. In this way, overall latency associated with process 400 may be reduced.

If a node change is not detected at step 404, then process 400 enters a sleep mode at step 406. For example, in some embodiments, an application thread or process may continuously check to determine if at least one node or link has changed in the network community. In other embodiments, the application thread or process may periodically check for changed links and nodes every n seconds, where n is any positive number. After the paths are calculated that go through a changed node at step 416 or after a period of sleep at step 406, process 400 may determine whether or not to loop at step 408. For example, if all changed nodes have been updated, then process 400 may stop at step 418. If, however, there are more changed nodes or links to process, then process 400 may loop at step 408 and return to step 404.

In practice, one or more steps shown in process 400 may be combined with other steps, performed in any suitable order, performed in parallel (e.g., simultaneously or substantially simultaneously), or removed.

FIGS. 4B-4E each include processes with a “map” phase and “reduce” phase. As described above, these phases may form part of a map/reduce computational paradigm carried out by parallel computational framework 114 (FIG. 1), key-value store 112 (FIG. 1), or both. As shown in FIG. 4B, in order to prepare any changed nodes, map phase 420 may include determining if there are any more link changes at step 422, retrieving the next link change at step 440, mapping the tail to out-link change at step 442, and mapping the head to in-link change at step 444.

If there are no more link changes at step 422, then, in reduce phase 424, a determination may be made at step 426 that there are more nodes and link changes to process. If so, then the next node and its link changes may be retrieved at step 428. The most recent link changes may be preserved at step 430 while any intermediate link changes are replaced by more recent changes. For example, the timestamp stored in table 306 (FIG. 3) may be used to determine the time of every link or node change. At step 432, the average out-link user connectivity value may be calculated. For example, if node n_1 has eight out-links with assigned user connectivity values, these eight user connectivity values may be averaged at step 432. At step 434, each out-link’s weight may be calculated in accordance with equation (1) above. All the out-link weights may then be summed and used to normalize each out-link weight at step 436. For example, each out-link weight may be divided by the sum of all out-link weights. This may yield a weight between 0 and 1 for each out-link. At step 438, the existing buckets for the changed node, in-links, and out-links may be saved. For example, the buckets may be saved in key-value store 112 (FIG. 1) or data store 110 (FIG. 1). If there are no more nodes and link changes to process at step 426, the process may stop at step 446.

As shown in FIG. 4C, in order to calculate paths originating from changed nodes, map phase 448 may include determining if there are any more changed nodes at step 450, retrieving the next changed node at step 466, marking existing buckets for deletion by mapping changed nodes to

12

the NULL path at step 468, recursively generating paths by following out-links at step 470, and if the path is a qualified path, mapping the tail to the path. Qualified paths may include paths that satisfy one or more predefined threshold functions. For example, a threshold function may specify a minimum path weight. Paths with path weights greater than the minimum path weight may be designated as qualified paths.

If there are no more changed nodes at step 450, then, in reduce phase 452, a determination may be made at step 454 that there are more nodes and paths to process. If so, then the next node and its paths may be retrieved at step 456. At step 458, buckets may be created by grouping paths by their head. If a bucket contains only the NULL path at step 460, then the corresponding cell in the node table may be deleted at step 462. If the bucket contains more than the NULL path, then at step 464 the bucket is saved to the corresponding cell in the node table. If there are no more nodes and paths to process at step 456, the process may stop at step 474.

As shown in FIG. 4D, in order to remove paths that go through a changed node, map phase 476 may include determining if there are any more changed nodes at step 478 and retrieving the next changed node at step 488. At step 490, the “bucket:” column in the node table (e.g., column 322 of node table 312 (both of FIG. 3B)) corresponding to the changed node may be scanned. For example, as described above, the target node identifier may be appended to the end of the “bucket:” column name. Each bucket may include a list of paths that connect the current node to the target node (e.g., the changed node). At step 492, for each matching node found by the scan and the changed node’s old buckets, the matching node may be matched to a (changed node, old bucket) deletion pair.

If there are no more changed nodes at step 478, then, in reduce phase 480, a determination may be made at step 484 that there are more node and deletion pairs to process. If so, then the next node and its deletion pairs may be retrieved at step 484. At step 486, for each deletion pair, any paths that go through the changed node in the old bucket may be deleted. If there are no more nodes and deletion pairs to process at step 482, the process may stop at step 494.

As shown in FIG. 4E, in order to calculate paths that go through a changed node, map phase 496 may include determining if there are any more changed nodes at step 498 and retrieving the next changed node at step 508. At step 510, the “bucket:” column in the node table (e.g., column 322 of node table 312 (both of FIG. 3B)) corresponding to the changed node may be scanned. At step 512, for each matching node found in the scan and the changed node’s paths, all paths in the scanned bucket may be joined with all paths of the changed bucket. At step 514, each matching node may be mapped to each qualified joined

If there are no more changed nodes at step 498, then, in reduce phase 500, a determination may be made at step 502 that there are more node and paths to process. If so, then the next node and its paths may be retrieved at step 504. Each path may then be added to the appropriate node bucket at step 506. If there are no more nodes and paths to process at step 502, the process may stop at step 516.

FIG. 5 shows illustrative process 520 for supporting a user query for all paths from a first node to a target node. For example, a first node (representing, for example, a first individual or entity) may wish to know how connected the first node is to some second node (representing, for example, a second individual or entity) in the network community. In the context of trust described above (and where the user connectivity values represent, for example, at least partially

13

subjective user trust values), this query may return an indication of how much the first node may trust the second node. In general, the more paths connecting the two nodes may yield a greater (or lesser if, for example, adverse ratings are used) network connectivity value (or network trust amount).

At step 522, the node table cell where the row identifier equals the first node identifier and the column equals the target node identifier appended to the “bucket:” column name prefix is accessed. All paths may be read from this cell at step 524. The path weights assigned to the paths read at step 524 may then be summed at step 526. At step 528, the path weights may be normalized by dividing each path weight by the computed sum of the path weights. A network connectivity value may then be computed at step 530. For example, each path’s user connectivity value may be multiplied by its normalized path weight. The network connectivity value may then be computed in some embodiments in accordance with:

$$t_{network} = \sum t_{path} \times w_{path} \quad (4)$$

where t_{path} is the user connectivity value for a path (given in accordance with equation (3)) and w_{path} is the normalized weight for that path. The network connectivity value may then be held or outputted (e.g., displayed on a display device, output by processing circuitry of application server 106, and/or stored on data store 110 (FIG. 1)). In addition, a decision-making algorithm may access the network connectivity value in order to make automatic decisions (e.g., automatic network-based decisions, such as authentication or identity requests) on behalf of the user. Network connectivity values may additionally or alternatively be outputted to external systems and processes located at third-parties. The external systems and processes may be configured to automatically initiate a transaction (or take some particular course of action) based, at least in part, on the received network connectivity values. Process 520 may stop at step 532.

In practice, one or more steps shown in process 520 may be combined with other steps, performed in any suitable order, performed in parallel (e.g., simultaneously or substantially simultaneously), or removed. In addition, as described above, various threshold functions may be used in order to reduce computational complexity. For example, a threshold function defining the maximum number of links to traverse may be defined. Paths containing more than the threshold specified by the threshold function may not be considered in the network connectivity determination. In addition, various threshold functions relating to link and path weights may be defined. Links or paths below the threshold weight specified by the threshold function may not be considered in the network connectivity determination.

Although process 520 describes a single user query for all paths from a first node to a target node, in actual implementations groups of nodes may initiate a single query for all the paths from each node in the group to a particular target node. For example, multiple members of a network community may all initiate a group query to a target node. Process 520 may return an individual network connectivity value for each querying node in the group or a single composite network connectivity value taking into account all the nodes in the querying group. For example, the individual network connectivity values may be averaged to form a composite value or some weighted average may be used. The weights assigned to each individual network connectivity value may be based on, for example, seniority in the community (e.g., how long each node has been a member in the community),

14

rank, or social stature. In addition, in some embodiments, a user may initiate a request for network connectivity values for multiple target nodes in a single query. For example, node n_1 may wish to determine network connectivity values between it and multiple other nodes. For example, the multiple other nodes may represent several candidates for initiating a particular transaction with node n_1 . By querying for all the network connectivity values in a single query, the computations may be distributed in a parallel fashion to multiple cores so that some or all of the results are computed substantially simultaneously.

In addition, queries may be initiated in a number of ways. For example, a user (represented by a source node) may identify another user (represented by a target node) in order to automatically initiate process 520. A user may identify the target node in any suitable way, for example, by selecting the target node from a visual display, graph, or tree, by inputting or selecting a username, handle, network address, email address, telephone number, geographic coordinates, or unique identifier associated with the target node, or by speaking a predetermined command (e.g., “query node 1” or “query node group 1, 5, 9” where 1, 5, and 9 represent unique node identifiers). After an identification of the target node or nodes is received, process 520 may be automatically executed. The results of the process (e.g., the individual or composite network connectivity values) may then be automatically sent to one or more third-party services or processes as described above.

In an embodiment, a user may utilize access application 102 to generate a user query that is sent to access application server 106 over communications network 104 (see also, FIG. 1) and automatically initiate process 520. For example, a user may access an Apple iOS, Android, or WebOS application or any suitable application for use in accessing application 106 over communications network 104. The application may display a searchable list of relationship data related to that user (e.g., “friend” or “follower” data) from one or more of Facebook, MySpace, openSocial, Friendster, Bebo, hi5, Orkut, PerfSpot, Yahoo! 360, LinkedIn, Twitter, Google Buzz, Really Simple Syndication readers or any other social networking website or information service. In some embodiments, a user may search for relationship data that is not readily listed—i.e., search Facebook, Twitter, or any suitable database of information for target nodes that are not displayed in the searchable list of relationship data. A user may select a target node as described above (e.g., select an item from a list of usernames representing a “friend” or “follower”) to request a measure of how connected the user is to the target node. Using the processes described with respect to FIGS. 3 and 4A-D, this query may return an indication of how much the user may trust the target node. The returned indication may be displayed to the user using any suitable indicator. In some embodiments, indicator may be a percentage that indicates how trustworthy the target node is to the user.

In some embodiments, a user may utilize access application 102 to provide manual assignments of at least partially subjective indications of how trustworthy the target node is. For example, the user may specify that he or she trusts a selected target node (e.g., a selected “friend” or “follower”) to a particular degree. The particular degree may be in the form of a percentage that represents the user’s perception of how trustworthy the target node is. The user may provide this indication before, after, or during process 520 described above. The indication provided by the user (e.g., the at least partially subjective indications of trustworthiness) may then be automatically sent to one or more third-party services or

15

processes as described above. In some embodiments, the indications provided by the user may cause a node and/or link to change in a network community. This change may cause a determination to be made that at least one node and/or link has changed in the network community, which in turn triggers various processes as described with respect to FIGS. 3 and 4A-4D.

In some embodiments, a path counting approach may be used in addition to or in place of the weighted link approach described above. Processing circuitry (e.g., of application server 106) may be configured to count the number of paths between a first node n_1 and a second node n_2 within a network community. A connectivity rating $R_{n_1n_2}$ may then be assigned to the nodes. The assigned connectivity rating may be proportional to the number of paths, or relationships, connecting the two nodes. A path with one or more intermediate nodes between the first node n_1 and the second node n_2 may be scaled by an appropriate number (e.g., the number of intermediate nodes) and this scaled number may be used to calculate the connectivity rating.

Each equation presented above should be construed as a class of equations of a similar kind, with the actual equation presented being one representative example of the class. For example, the equations presented above include all mathematically equivalent versions of those equations, reductions, simplifications, normalizations, and other equations of the same degree.

The above described embodiments of the invention are presented for purposes of illustration and not of limitation. The following numbered paragraphs give additional embodiments of the present invention.

What is claimed is:

1. A method for determining network connectivity between a first node and a second node connected to the first node by at least one path, the method comprising:

identifying paths to the second node from the first node within a representation of a network community stored in one or more electronic databases, wherein each path comprises one or more links, and wherein each link is assigned a user connectivity value;

using processing circuitry to:

automatically, without user input, identify sub-processes required to calculate a network connectivity value, wherein identifying the sub-processes comprises:

identifying a plurality of links in the identified paths; for each identified link, accessing a data structure to identify nodes connected to the identified link; and for each identified node, creating an indication of a sub-process, wherein the sub-process comprises calculating a normalized out-link weight for each out-link of the identified node;

distribute the indications of the sub-processes to a plurality of processors arranged in a parallel computational framework;

receive, from the plurality of processors, the calculated normalized out-link weights for the identified nodes; determine a normalized path weight based on a plurality of the calculated normalized out-link weights for each of the identified paths;

for each of the identified paths, identify a minimum user connectivity value assigned to a link in the identified path;

for each of the identified paths, calculate a product of the normalized path weight for the identified path and the identified minimum user connectivity value assigned to the link in the identified path;

16

sum the calculated products to produce a network connectivity indication;

output the network connectivity indication;

determine whether one or more user connectivity values for an identified path has changed; and

calculate at least one prospective decision based on the determination.

2. The method of claim 1 further comprising accessing a link threshold value, wherein identifying the paths to the second node from the first node within the representation of the network community comprises identifying only those paths containing fewer links than the accessed link threshold value.

3. The method of claim 1 further comprising accessing a path weight threshold value, wherein identifying the paths to the second node from the first node within the representation of the network community comprises identifying only those paths with a normalized path weight above the accessed path weight threshold value.

4. The method of claim 1 wherein the processing circuitry determines the normalized path weight for one of the identified paths by dividing a product of the calculated normalized out-link weights of links in the one path by a sum of path weights of all the identified paths.

5. The method of claim 1 wherein the user connectivity value represents at least one of a subjective user trust value or a competency assessment.

6. The method of claim 5 further comprising using the processing circuitry to

access third-party ratings data; and

determine the network connectivity indication based, at least in part, on the third-party ratings data.

7. The method of claim 1 wherein identifying the paths to the second node from the first node within the representation of the network community comprises accessing data from a social networking service.

8. The method of claim 1 wherein identifying the paths to the second node from the first node within the representation of the network community comprises retrieving a pre-stored identification of the paths to the second node from the first node from a table in a database.

9. The method of claim 1 further comprising automatically making at least one network-based decision based, at least in part, on the network connectivity indication.

10. The method of claim 1 wherein the processing circuitry is further configured to automatically determine that at least one node in the network community has changed and in response to the determination, update the network connectivity indication.

11. A system for determining network connectivity between a first node and a second node connected to the first node by at least one path, the system comprising processing circuitry configured to:

identify paths to the second node from the first node within a representation of a network community stored in one or more electronic databases, wherein each path comprises one or more links, and wherein each link is assigned a user connectivity value;

automatically, without user input, identify sub-processes required to calculate a network connectivity value, wherein identifying the sub-processes comprises:

identifying a plurality of links in the identified paths; for each identified link, accessing a data structure to identify nodes connected to the identified link; and for each of the identified nodes, creating an indication of a sub-process, wherein the sub-process comprises

17

calculating a normalized out-link weight for each out-link of the identified node;
 distribute the indications of the sub-processes to a plurality of processors arranged in a parallel computational framework;
 receive, from the plurality of processors, the calculated normalized out-link weights for the identified nodes;
 determine a normalized path weight based on a plurality of the calculated normalized out-link weights for each of the identified paths;
 for each of the identified paths, identify a minimum user connectivity value assigned to a link in the identified path;
 for each of the identified paths, calculate a product of the normalized path weight for the identified path and the identified minimum user connectivity value assigned to the link in the identified path;
 sum the calculated products to produce a network connectivity indication;
 output the network connectivity indication;
 determine whether one or more user connectivity values for an identified path has changed; and
 calculate at least one prospective decision based on the determination.

12. The system of claim 11 wherein the processing circuitry is further configured to access a link threshold value, wherein the processing circuitry identifies the paths to the second node from the first node within the representation of the network community by identifying only those paths containing fewer links than the accessed link threshold value.

13. The system of claim 11 wherein the processing circuitry is further configured to access a path weight threshold value, wherein the processing circuitry identifies the paths to the second node from the first node within the

18

representation of the network community by identifying only those paths with a normalized path weight above the accessed path weight threshold value.

14. The system of claim 11 wherein the processing circuitry determines the normalized path weight for one of the identified paths by dividing a product of the calculated normalized out-link weights of links in the one path by a sum of path weights of all the identified paths.

15. The system of claim 11 wherein the user connectivity value represents at least one of a subjective user trust value or a competency assessment.

16. The system of claim 15 wherein the processing circuitry is further configured to

access third-party ratings data; and

determine the network connectivity indication based, at least in part, on the third-party ratings data.

17. The system of claim 11 wherein the processing circuitry identifies the paths to the second node from the first node within the representation of the network community by accessing data from a social networking service.

18. The system of claim 11 wherein the processing circuitry identifies the paths to the second node from the first node within the representation of the network community by retrieving a pre-stored identification of the paths to the second node from the first node from a table in a database.

19. The system of claim 11 wherein the processing circuitry is further configured to automatically make at least one network-based decision based, at least in part, on the network connectivity indication.

20. The system of claim 11 wherein the processing circuitry is further configured to automatically determine that at least one node in the network community has changed and in response to the determination, update the network connectivity indication.

* * * * *